arXiv:1612.01810v1 [cs.CV] 6 Dec 2016

FLIC: Fast Linear Iterative Clustering with Active Search

Jia-Xin Zhao Ren Bo Qibin Hou Ming-Ming Cheng
CCCE, Nankai University
http://mmcheng.net/flic/

Abstract

Benefiting from its high efficiency and simplicity, Simple Linear Iterative Clustering (SLIC) remains one of the most popular over-segmentation tools. However, due to explicit enforcement of spatial similarity for region continuity, the boundary adaptation of SLIC is sub-optimal. It also has drawbacks on convergence rate as a result of both the fixed search region and separately doing the assignment step and the update step. In this paper, we propose an alternative approach to fix the inherent limitations of SLIC. In our approach, each pixel actively searches its corresponding segment under the help of its neighboring pixels, which naturally enables region coherence without being harmful to boundary adaptation. We also jointly perform the assignment and update steps, allowing high convergence rate. Extensive evaluations on Berkeley segmentation benchmark verify that our method outperforms competitive methods under various evaluation metrics. It also has the lowest time cost among existing methods (approximately $30fps$ *for a* 481×321 *image on a single CPU core).*

1. Introduction

Superpixels, generated by image over-segmentation, take the place of pixels to become the fundamental units in various computer vision tasks, including image segmentation [18], image classification [26], 3D reconstruction [9], object tracking [25], *etc*. Such a technique can greatly reduce the computation complexity, avoid under-segmentation, and reduce the influence caused by noise. Moreover, it preserves useful information as much as possible, and hence fits the human visual system better. Therefore, how to generate superpixels with high efficiency plays an important role in most vision and image processing applications.

Generating superpixels has been an important research issue, with a group of classical methods been developed, including FH [8], Mean Shift [6], Watershed [23], *etc*. Though these methods have been widely used recently, they suffer from lack of compactness and produce irregular superpixels especially when contrast is poor or shadow exists. To solve the above-mentioned problems, Shi and Malik proposed Normalized Cuts (NC) [19] that generated compact superpixels. However, this method is hard to adhere to image boundaries very well, and the high complexity limits its application. GraphCut [5, 22] was proposed regarding the segmentation problem as an energy optimization process. It solved the compactness problem by using min-cut/max-flow algorithms [4, 10] in GraphCut, but too many parameters made it much harder to be used. Turbopixel [11] is another method that is proposed to solve the compactness problem. However, the inefficiency of the underlying level-set method [16] restricts its applications. Bergh *et al.* proposed an energy-driven algorithm SEEDS [21] whose results adhered to the boundaries well, but unfortunately it suffers from irregularity and the number of superpixels is difficult to control. ERS [12], although performs well on the Berkeley segmentation benchmark, has the high computational cost that limits its practical use.

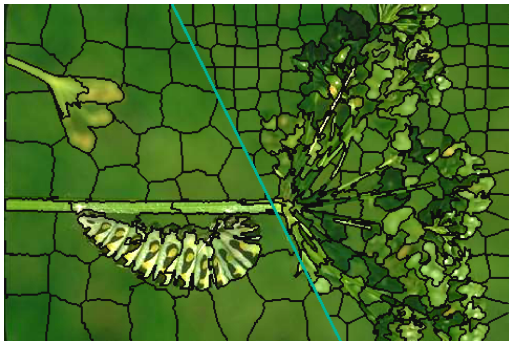

Figure 1. Segmentation results of our proposed approach with 100 (left) and 400 (right) superpixels.

Achanta *et al.* proposed a linear clustering based algorithm SLIC [1]. It generates superpixels based on the Lloyd's algorithm [14] that is also known as Voronoi iteration or k-means. In the assignment step of SLIC, as a key point to speed up the algorithm, each pixel p is associated with those cluster seeds whose search regions overlap its location. Such a strategy is also adopted by most subsequent works based on SLIC. SLIC is widely used in

various applications [25] because of its high efficiency and good performance. Inspired by SLIC, Wang *et al.* [24] implemented an algorithm SSS that considered the structure information of images. It uses the geodesic distance [17] computed by the geometric flows instead of the simple Euclidean distance. However, efficiency becomes its bottleneck due to the high computational cost of measuring the geodesic distance. Very recently, Liu *et al.* proposed Manifold SLIC [13] that generated content-sensitive superpixels by computing Centroidal Voronoi Tessellation (CVT) [7] in a special feature space. Such an advanced technique makes it much faster than SSS but still slower than SLIC owning to its cost of mapping, splitting and merging processes. From the aforementioned description, we know that the descendants of SLIC improve the results by either using more complicated distance measurements or providing more suitable transformations of the feature space. However, the assignment and update steps of these methods are performed separately and are not changed too much, so they cannot converge well until many times of iteration.

In this paper, we consider the over-segmentation problem from a new perspective. Each pixel, in our algorithm, is allowed to actively search which superpixel it should belong to according to its neighboring pixels as shown in Fig. 2. At the meantime, the seeds of the superpixels can be adaptively changed during this process, which allows our assignment and update steps to be performed jointly. This property enable our approach to reach convergence at a very fast speed. To sum up, our main advantages can be concluded as follows:

- Our algorithm features well-awareness of neighboring-pixel continuity and provides results of good boundary sensitivity regardless of image complexity and contrast.
- Our algorithm allows to perform the assignment step and the update step in a jointed manner, and has a high convergence rate as well as the lowest time cost among all superpixels approaches. Experiments show that our approach is able to converge in two scan loops, with better performance measured under a variety of evaluation metrics on the Berkeley segmentation benchmark.

2. Preliminaries

SLIC improves Lloyd's algorithm, reducing the time complexity from $O(KN)$ to $O(N)$, where K is the number of the superpixels and N is the number of pixels. Due to its high efficiency and simplicity, SLIC has been widely used in a variety of vision-related tasks. Since our proposed approach is closely related to SLIC, we will briefly recap it and its followups in this section.

Let $\{I_i\}_{i=1}^{N}$ be a colorful image, where I_i represents the corresponding variable of each pixel. Given a set of evenly distributed seeds $\{S_k\}_{k=1}^{K}$, SLIC simplifies the Lloyd's algorithm to get the Centroidal Voronoi Tessellation (CVT) [7] that will be introduced in Sec. 3.3. In the assignment step, each pixel I_i is associated with those cluster seeds whose search regions overlap its location as shown in Fig. 2(a). The area of a search region can be denoted by $2T \times 2T$, where $T = \sqrt{N/K}$. Specifically, SLIC denotes I_i in a five dimensional space that contains a three dimensional CIELAB color space (l_i, a_i, b_i) and a two dimensional spatial space (x_i, y_i). SLIC measures the distance between two points using the Euclidean distance, which can be computed by

$$D(I_i, I_j) = \sqrt{d_c^2 + \left(\frac{d_s * m}{N_s}\right)^2}, \tag{1}$$

where m is a variable that controls the weight of the spatial term, and $N_s = T$. Variables d_s and d_c are respectively the spatial and color distances, which can be expressed as

$$d_s = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}, \tag{2}$$

as well as

$$d_c = \sqrt{(l_i - l_j)^2 + (a_i - a_j)^2 + (b_i - b_j)^2}. \tag{3}$$

In the update step, SLIC recomputes the center of each superpixel and moves the seeds to these new centers. Then it obtains the over-segmentation results by iteratively performing the assignment and update steps.

The follow-up works of SLIC also use similar procedure as SLIC. They improve the performance of SLIC using better distance measures or more suitable transformation function between color space and spatial space. However, in SLIC-based algorithms, each search region is fixed in the assignment step of a single loop, and the region continuity information of neighboring pixels is largely ignored when allocating pixels to superpixels. Separately performing the assignment step and the update step also leads to a delayed feedback of pixel label change.

3. Fast Over-Segmentation with Active Search

Our algorithm is shown in Algorithm 1. Each pixel can actively choose which superpixel it should belong to in a forth-and-back order, and the assignment and update steps are performed jointly. We will explain them in detail in this section.

3.1. Problem Setup

Given the desired number of superpixels K and an input image $I = \{I_i\}_{i=1}^{N}$, where N is the number of pixels, our goal is to produce a series of disjoint small regions (or superpixels, we call them lumps in this paper) as output. As in most previous works [1], the original RGB color space is

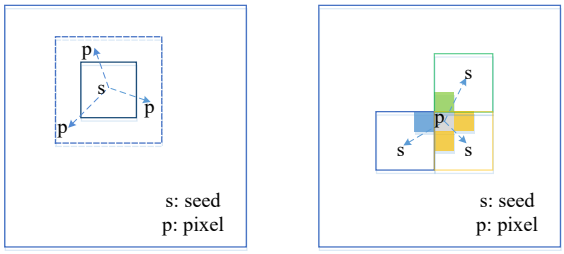

(a) SLIC (b) Active search

Figure 2. (a) The search method used in SLIC. Each seed only searches a limited region to reduce computation complexity. (b) Our proposed active search. Each pixel is able to decide its own label by searching its surroundings.

transformed to the CIELAB color space. Thus, each pixel I_i in an image I can be represented in a five dimensional space,

$$I_i = (l_i, a_i, b_i, x_i, y_i). \tag{4}$$

The original image is further divided into K regular grids $\{G_k\}_{k=1}^{K}$ with step length $\upsilon = \sqrt{N/K}$ and the initial label for each pixel I_i is assigned as:

$$l_i = k, \text{ if } I_i \in G_k. \tag{5}$$

We initialize each seed S_k as the mass center of G_k. Therefore, S_k can also be defined in the same five dimensional space

$$S_k = \{l_k, a_k, b_k, x_k, y_k\}. \tag{6}$$

After initialization, traditionally the Lloyd's algorithm (k-means) is used to find the global or local optimal solution. In this paper, we propose a novel algorithm to get the image over-segmentation.

3.2. Active search method

The difference between a natural image and a set of colored discrete points is the abundant priori knowledge, especially neighboring continuity, in the natural image. Considering local continuity information of pixels, we know that in most images, adjacent pixels tend to have the same labels, *i.e.*, neighboring pixels have natural continuity. Hence we let the pixels actively search for the nearest seeds of superpixels as shown in Fig. 2. In such an "active search", we utilize the priori local continuity information, assuming that neighboring pixels are more likely to have the same labels. We then only compute the distance between a pixel and the seeds of its four neighboring pixels. Specifically, for a pixel I_i, our assignment principle is

$$l_i = \underset{l_j}{argmin}\, D(I_i, S_{l_j}), I_j \in A_i, \tag{7}$$

where A_i consists of I_i and its four neighboring pixels, S_{l_j} is I_j's corresponding superpixel seed. We use Eqn. (1) to measure the distance $D(I_i, S_{l_j})$.

Algorithm 1 FLIC

Input: Image I with N pixels, the desired number of superpixels K, the maximal iteration numbers itr_{max} and the spatial distance weight m.
Output: K superpixels

```
Divide the original image into regular grids {G_k}_{k=1}^K
  with step length υ = sqrt(N/K).
Initialize labels {l_k}_{k=1}^K for pixels according to their
  locations.
Move each seed to the lowest gradient position in a 3×3
  neighborhoods.
Initialize seeds {S_k}_{k=1}^K.
Regard the pixels sharing the same label as a lump ζ.
Initialize distance d_(i) = ∞ for each pixel and itr = 0.
while itr < itr_max do
    for each lump ζ_k do
        Use forth-and-back scan to traverse the lump ζ_k
          to get the pixels processing sequence. (Sec. 3.2).
        for each pixel I_i in the sequence do
            Set d_(i) = D(I_i, S_{l_i}) by Eqn. (1)
            for I_j in the four-neighborhood of I_i do
                if l_j ≠ l_i then
                    Compute D = D(I_i, s_{l_{p_1}}).(Eqn. (1))
                    if D < d_(i) then
                        d_(i) = D; l_i = l_j.
                    end if
                end if
            end for
            if l_i is changed to l_j then
                Use Eqn. (10) to update ζ_{l_i};
                Use Eqn. (11) to update ζ_{l_j};
                Update the bounding box of ζ_{l_j}.(Sec. 3.3)
            end if
        end for
    end for
    itr++;
end while
```

Since pixels can only be assigned to a superpixel containing one of its neighbors, the local pixels continuity has a stronger effect in the proposed strategy, *i.e.*, the pixels actively assign themselves to one of the neighboring connected superpixel regions. Note that such assignment does not have a fixed range limit in space. As a result, Eqn. (7) leads to better boundary adherence. Detailed data and analysis will be provided in Sec. 4.3. During this process the superpixel centers are also self-adaptively modified for faster convergence, which we will discuss in detail in Sec. 4.3.

From Eqn. (7), the pixel traverse order can influence the assignment result since formerly traversed pixels may already experienced label-change. Fortunately, we can use a forth-and-back traverse order as in PatchMatch [3], in

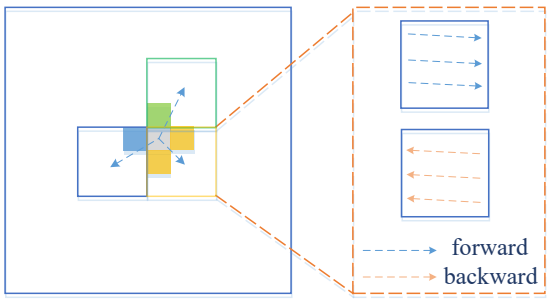

(a) Processing pixel (b) Processing lump

Figure 3. Illustration of processing order in our method. We process the lumps independently. The order of pixels processing for each lump is obtained by force-and- back scan method which contains a forward and a backward pass (b). Then each pixel in above sequence can be processed as shown in (a).

which the pixels later processed will benefit from the previous pixels. Specifically, within the bounding box of each lump, we first traverse the pixels in the lump in normal forward scan-line order, using Eqn. (7) to assign each pixel's label, then the traverse order is reversed, and Eqn. (7) is applied again to these pixels. During scan-line order processing, pixel assignments could only consider information from upper-left (previous) pixels. Thus, we process each lump in a forth-and-back order as shown in Fig. 3 so that information from all directions could be considered and our processing could be more adequate.

3.3. Jointed step for assignment and update

One problem in previous methods such as SLIC is that the assignment step and the update step are performed separately, leading to delayed feedbacks from pixel label change to superpixel seeds. As a result, multiple iterations are required before the algorithms converge.

In our approach, based on the assignment principle Eqn. (7), we design a jointed assignment and update strategy that is able to adjust superpixel seed center position on the run, drastically reducing the iteration number needed for convergence. Most clustering-based methods aim to get the Centroidal Voronoi tessellation (CVT). Hence we will briefly introduce the CVT.

Let $S = \{S_k\}_{k=1}^{K}$ be the set of seeds in the image, where K is the expected number of superpixels. The Voronoi cell $\mathcal{V}_{(S_k)}$ of a seed S_k is denoted by:

$$\mathcal{V}_{(S_k)} = \{I_i \in I | d(I_i, S_k) \leqslant d(I_i, (S_j)), \forall S_j \in S\}, \quad (8)$$

where $d(I_i, S_k)$ is the distance from pixel I_i to the seed S_k. The Voronoi Diagram $\mathcal{V}_D(S)$ is defined by

$$\mathcal{V}_D(S) = \{\mathcal{V}_{(S_k)} \neq \phi | \forall S_k \in S\}. \quad (9)$$

A CVT is then defined as a Voronoi Diagram whose generator point of each Voronoi cell is also its center of mass. As mentioned above, traditionally CVT is usually obtained by heuristic algorithms such as Lloyd's algorithm and iteratively updated after each assignment step until convergence is reached.

In our approach, using Eqn. (7) allows us to joint the update step and the assignment step together. After pixel I_i is processed, if its label is changed to l_j we immediately update the current seed S_{l_i} using the following equation:

$$S_{l_i} = \frac{S_{l_i} * |\zeta_{l_i}| - I_i}{|\zeta_{l_i}| - 1}, \quad (10)$$

where $|\zeta_{l_i}|$ is the number of pixels in the lump ζ_{l_i}, and using the following equation to update S_{l_j}

$$S_{l_j} = \frac{S_{l_j} * |\zeta_{l_j}| + I_i}{|\zeta_{l_j}| + 1}. \quad (11)$$

Besides, we also need to update the bounding box of ζ_{l_j}.

The above mentioned method only contains simple arithmetic operations and hence can be performed very efficiently. Such immediate update will help later pixels make a better choice during assignment, contributing largely to better convergence property.

4. Experiments

Our method is implemented by C++ and runs on a PC with an Intel Core i7-4790K CPU with 4.0GHz, 32GB RAM, and 64 bit operation system. We compare our method with many existing works, including FH [8], SLIC [1], Manifold SLIC [13], SEEDS [21], and ERS [12] on the BSDS500 benchmark with the evaluation method [2, 20]. In the test phase, we evaluate the algorithms on 200 images with resolution 481 $\times$ 321 from Berkeley dataset.

4.1. Parameters

In our approach, three parameters need to be set: (i) The number of superpixels K. The common advantage of the clustering-based algorithm is that the expected number superpixels are able to be got easily by set the parameter K; (ii) A spatial distance weight m. Parameter m makes a great effect on the superpixels smooth. We shall show that our performance will increase along with the decrease of m. However, a small m can also lead to the irregularity of superpixels. To achieve a great trade-off between the compactness and the performance, in the following experiments, we set $m = 5$ in default; (iii) The number of iterations itr. In Fig. 6(b), we show the Boundary Recall - itr curve to illustrate the high convergence rate of FLIC. Hence we set $itr = 2$ in default to get the balance between time cost and performance.

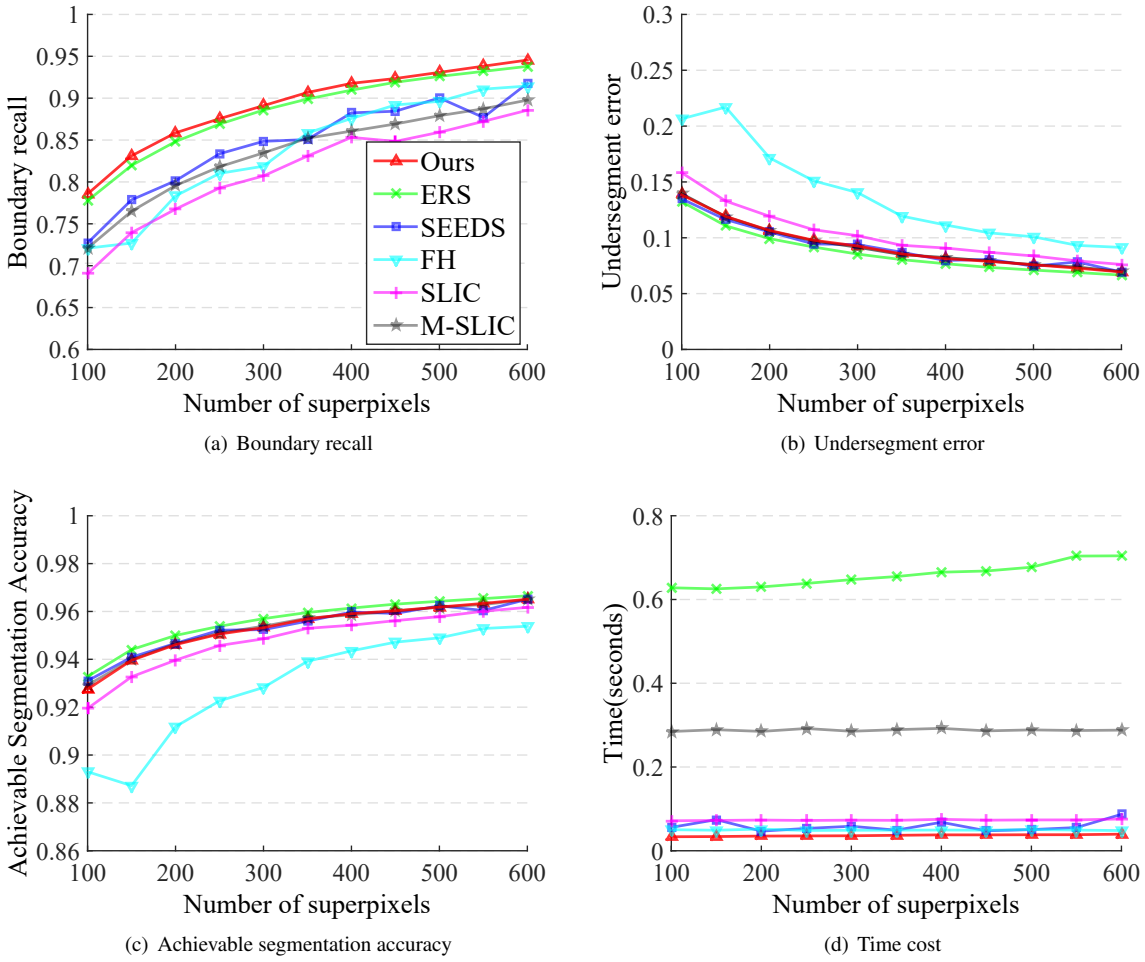

(a) Boundary recall

(b) Undersegment error

(c) Achievable segmentation accuracy

(d) Time cost

Figure 4. Evaluation of representative algorithms and our method on the BSDS500 benchmark for $K \in [100, 600]$. Over-segmentation generated by our method FLIC has the highest boundary recall ratio, reflecting that our method adheres to boundaries very well, and our result achieves the competitive result with the state-of-the-art method ERS [12] in achievable segmentation accuracy and undersegment error. However, FLIC run 20 times faster than ERS.

4.2. Compared to the state-of-the-arts

Boundary recall (BR). Superpixels are usually used to replace original pixels as the based units, so it should preserve the boundaries well. Boundary recall is a measurement which denotes the adherence to the boundaries. It computes what fraction of the ground truth edges falls within ε-pixel length from at least one superpixel boundary. The BR [1] can be computed by

$$BR_{\mathcal{G}}(\mathcal{S}) = \frac{\sum_{p \in \xi_{\mathcal{G}}} \Pi(\min_{q \in \xi_{\mathcal{S}}} \|p - q\| < \varepsilon)}{|\xi_{\mathcal{G}}|}. \quad (12)$$

, where $\xi_{\mathcal{S}}$ and $\xi_{\mathcal{G}}$ denote the union set of superpixel boundary and the union set of ground truth boundaries, respectively. The indicator function Π checks if the nearest pixel is within ε distance. Here we follow previous art [1, 13] to set $\varepsilon = 2$ in our experiment.

The boundary recall curves of different methods are plotted in Fig. 4(a). One can observe that our FLIC method outperforms all other methods. When $K = 200$, the recall rate of our result is 0.859, which is significantly higher than 0.796 in Manifold SLIC [13] and slightly better than the ERS [12] method.

Undersegment error (UE). The undersegment error reflects the extent that superpixels do not exactly overlap the ground truth segmentation. Similar to boundary recall, UE can also reflect the boundary adherence. The difference is that UE uses segmentation regions instead of boundaries in measurement. Mathematically, the UE [15] can be computed by

$$UE_{\mathcal{G}}(\mathcal{S}) = \frac{\sum_{G \in \varrho_{\mathcal{G}}} (\sum_{S:S \cap G \neq \phi} \min(S_{in}, S_{out}))}{N}, \quad (13)$$

where $\varrho_{\mathcal{S}}$ is the union set of superpixels, $\varrho_{\mathcal{G}}$ is the union set

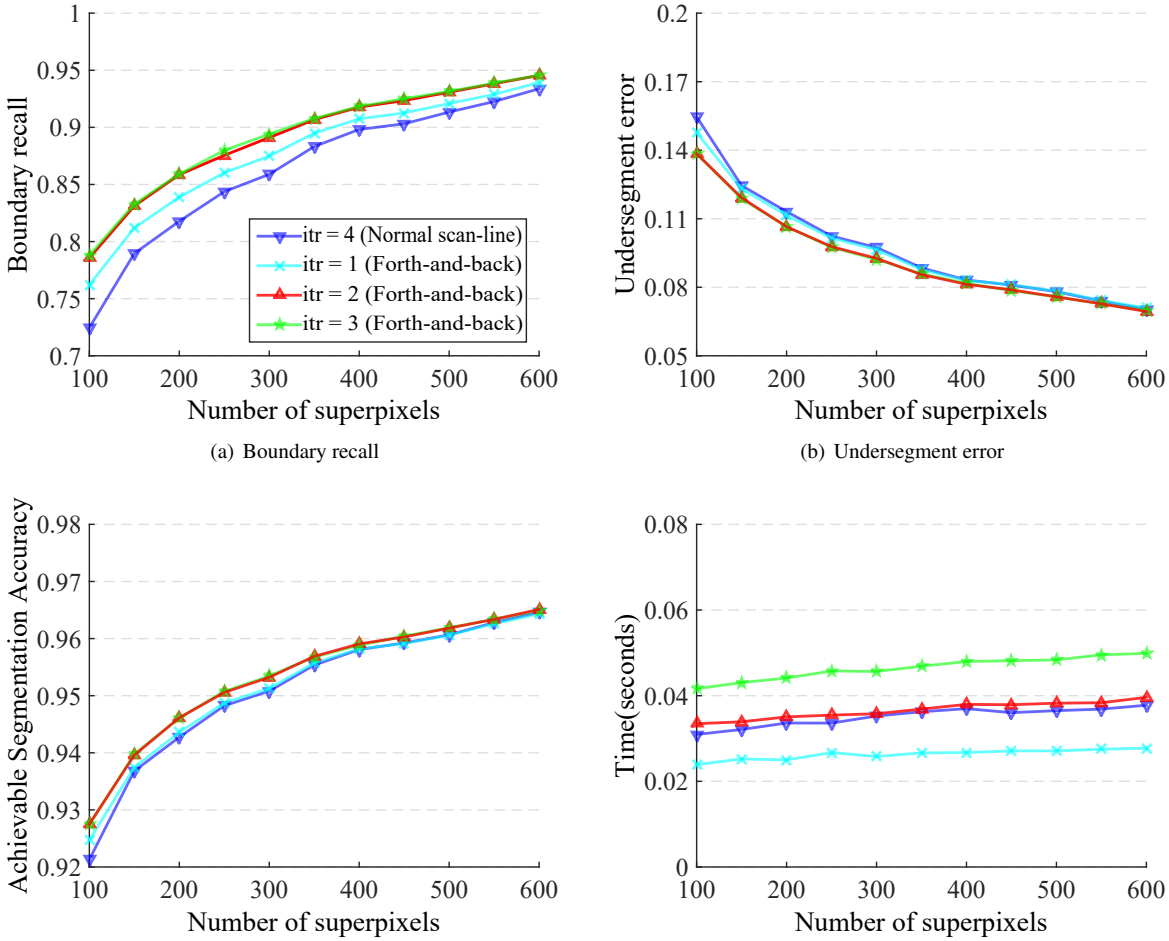

(a) Boundary recall

(b) Undersegment error

(c) Achievable segmentation accuracy

(d) Time cost

Figure 5. Part of ablation analysis under the standard evaluation metrics and time cost.

of the segments of the ground truth, S_{in} denotes the overlapping of the superpixel S and the ground truth segment G, and S_{out} denotes the rest of the superpixel S.

As shown in Fig. 4(b), our method produces very close results to previous state-of-the-arts. The undersegment error of the superpixels generated by FLIC is 0.108 when K is set to 200, while the UE value of the results produced by the state-of-the-art method ERS [12] with the same superpixel number is 0.10.

Achievable segmentation accuracy (ASA). ASA gives the highest accuracy achievable for object segmentation that utilizes superpixels as units. Similar to undersegment error, ASA utilizes segments instead of the boundaries, which can be computed by [12]

$$ASA_{\mathcal{G}}(\mathcal{S}) = \frac{\sum_k \max_i |S_k \cap G_i|}{\sum_i G_i}, \tag{14}$$

where S_k represents superpixel and G_i represents ground truth segment. Better segmentation of superpixels will have a larger ASA value.

As shown in Fig. 4(c), compared to the state-of-the-art [12], the performance of our approach is competetive. For example, the ASA value of the superpixels generated by FLIC with 200 superpixels is 0.946. The results with the same number of superpixels is 0.95 with the state-of-the-art method ERS [12].

Time cost (TC). Similar to SLIC, our method has an $O(N)$ time complexity, which is independent of number of superpixels K. It is crucial for many practical tasks since superpixel generation is pre-requisite of many applications. In fact, time cost is the major superiority of the SLIC[1] algorithm, and many approaches are limited by their speeds such as SSS[24] and ERS [12], *etc*.

As shown in Fig. 4(d), the average time cost of FLIC with two iterations is 0.035s, while the time cost of ERS, Manifold SLIC, SLIC and FH is 0.625, 0.281s, 0.072s, and 0.047s, respectively. FLIC has the lowest time cost among all methods. It runs nearly 20 times faster than the ERS

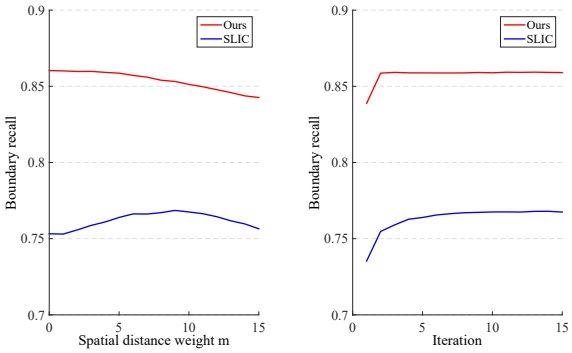

(a) (b)

Figure 6. (a) The boundary recall - m curves, where m is the spatial distance weight in Eqn. (1). Our overall performance is far better for all the tested m. (b) The boundary recall - Iteration curves. Our method converges within 2 iterations, which is much faster than SLIC.

method with comparable result quality.

4.3. Ablation analysis

Effect of the forth-and-back traverse order. In Fig. 5, we compare the performance with and without such strategy. Specifically, we compare the results using normal scan method to traverse all lumps four times (blue) with that using the proposed scan method two times (red). From Fig. 5, we can observe that the red curve significantly outperforms blue curve under all standard evaluation metrics while achieving competitive time cost compared to blue curve.

Effect of the spatial distance weight. As shown in Fig. 6(a), on the contrary to the SLIC method, the boundary recall curve with respect to the spatial distance weight m is monotonously decreasing. The reason for this phenomenon is that in our method local region continuity is mostly ensured by the active search algorithm, and color boundary are less preserved for bigger m value. On the other hand, a too small m will result in less regularity of superpixels, so we choose $m = 5$ in our comparison with previous works. Nonetheless, our overall performance is significantly better for all m values.

Convergence. FLIC significantly accelerates the evolution so that we only need a few iterations before convergence. We compare the performance curves with different iterations on Berkeley benchmark in Fig. 5. Our algorithm quickly converges within two iterations, more iterations only bring marginal benefits to the result. Numerically, the boundary recall of the superpixels with only one iteration is 0.835 when K is set to 200.

The value of two iterations is 0.859 and three iterations is 0.86 when generating the same number of superpixels. The undersegment error values are 0.115, 0.108, and 0.107, respectively. The achievable segmentation accuracy values are 0.941, 0.945, and 0.946, respectively. As can be seen in Fig. 6(b), our algorithm not only converges much faster than SLIC (which costs 10 iterations to converge), but also obtains a better performance in the Boundary Recall measurement.

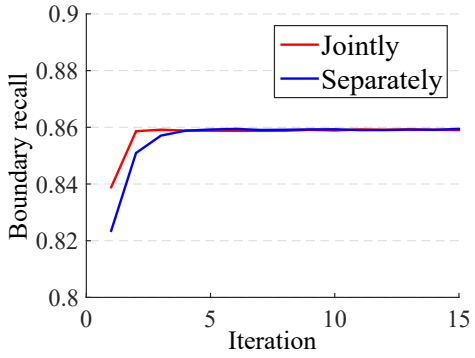

Figure 7. Comparison of the convergence rate between jointed and separated assignment and update steps.

Role of the jointed assignment and update. Our algorithm allows jointly performing the assignment and update steps. In Fig. 7 we show the convergence rate of both our jointed approach and that of traditionally separating the assignment and update steps in our calculation. The jointed approach converges within two iterations, while the iteration number needed for convergence doubles for the separate calculation, which means a 50% performance enhancement.

Effect on size of neighborhood. In our implementation, we use four neighboring pixels for pixel assignment. It is possible to search more pixels, say 8 neighboring pixels, in this process. In the Table 1 and Table 2, we briefly compare the effect of different sizes of neighborhood. A larger neighborhood brings better performance, but sacrifices running speed.

Boundary Recall(%)				
	100	200	300	400
4-neighborhood	78.6	85.9	89.1	91.8
8-neighborhood	80.5	87.4	90.5	92.7

Table 1. Boundary recall comparison between 4-neighborhood and 8-neighborhood with different superpixel count 100, 200, 300, 400.

Time cost(ms/image)				
	100	200	300	400
4-neighborhood	34	35	36	38
8-neighborhood	54	56	59	61

Table 2. Time cost comparison between 4-neighborhood and 8-neighborhood with different superpixel count 100, 200, 300, 400.

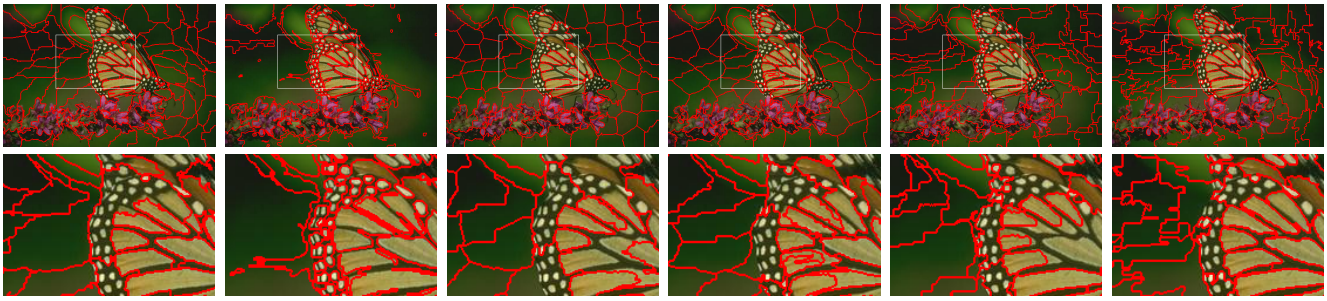

(a) Ours (b) FH [8] (c) SLIC [1] (d) M-SLIC [13] (e) SEEDS [21] (f) ERS [12]

Figure 8. Visual comparison of superpixel segmentation results using different existing algorithm with 100 superpixels.

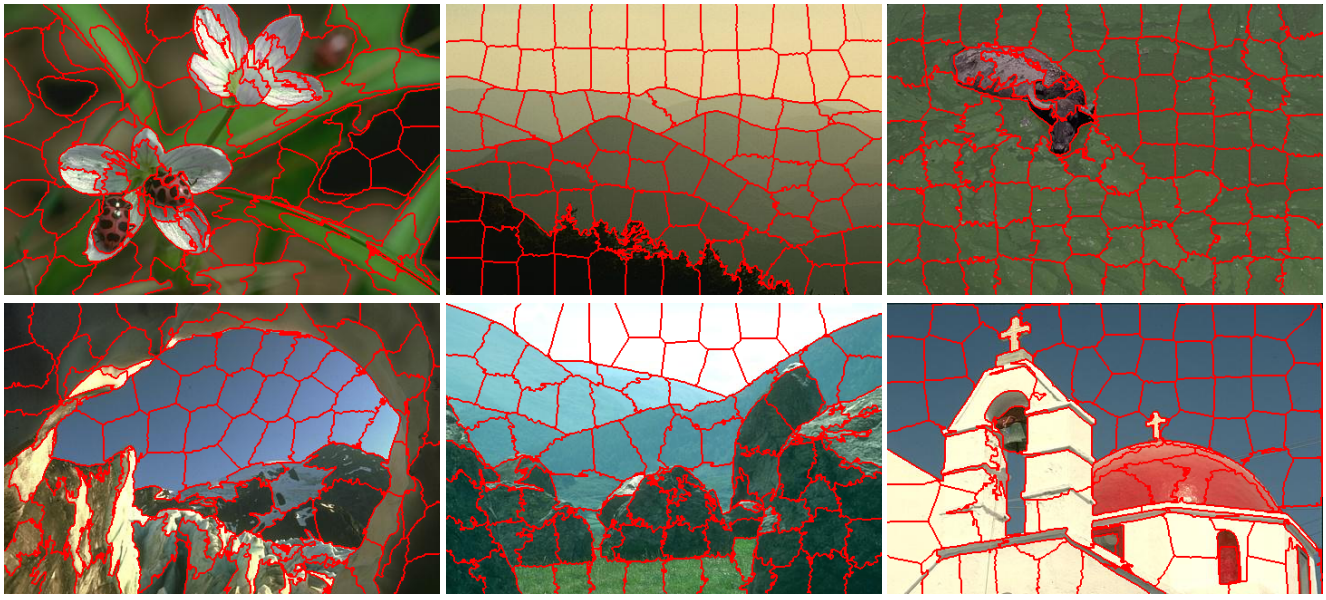

Figure 9. Superpixel over-segmentation examples, and each image contains 100 superpixels.

4.4. Qualitative Results

In Fig. 8, we list several superpixel segmentation results using different algorithms. Fig. 8(a) shows our approach; (b) shows FH [8], which is a graph-based algorithm; (c) shows SLIC [1], which achieves an compromise between accuracy and efficiency but performs mediocrity under the standard evaluation metrics; (d) shows Manifold SLIC [13], which is content-sensitive but slow for practical applications; (e) shows SEEDS [21], which run faster than most approaches but is less compact and difficult to control the superpixel count; (f) shows ERS [12], which performs very well under all the standard evaluation metrics but runs too slow to use in practical and suffer from irregularity problems.

We are able to conclude that our approach is more sensitive to image boundary, especially under deficient contrast between background and object, and achieves an excellent compromise between adherence and compactness. More examples of our over-segmentation results are provided in Fig. 9 and Fig. 1.

5. Conclusions

In this paper we present a novel algorithm using an active search, which is able to improve the performance and significantly reduce the time cost for over-segmentation. Taking advantage of local continuity, our algorithm provides results of good boundary sensitivity regardless of image contrast and complexity. Moreover, it is able to converge in only two iterations, achieving the lowest time cost compared to previous methods while obtaining performance comparable to the state-of-the-art method ERS with 1/20 running time. We have used various evaluation metrics on the Berkeley segmentation benchmark to demonstrate the high efficiency and high performance of our approach. In the future, we plan to improve our results by exploring more advanced distance measurements, and further extend our method to deal with 3D video data.

References

[1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE TPAMI*, 34(11):2274–2282, 2012. 1, 2, 4, 5, 6, 8

[2] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2011. 4

[3] C. Barnes, E. Shechtman, A. Finkelstein, and D. Goldman. Patchmatch: a randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics-TOG*, 28(3):24, 2009. 3

[4] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE transactions on pattern analysis and machine intelligence*, 26(9):1124–1137, 2004. 1

[5] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on pattern analysis and machine intelligence*, 23(11):1222–1239, 2001. 1

[6] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603–619, 2002. 1

[7] Q. Du, V. Faber, and M. Gunzburger. Centroidal voronoi tessellations: applications and algorithms. *SIAM review*, 41(4):637–676, 1999. 2

[8] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59(2):167–181, 2004. 1, 4, 8

[9] D. Hoiem, A. A. Efros, and M. Hebert. Automatic photo pop-up. *ACM transactions on graphics (TOG)*, 24(3):577–584, 2005. 1

[10] V. Kolmogorov and R. Zabin. What energy functions can be minimized via graph cuts? *IEEE transactions on pattern analysis and machine intelligence*, 26(2):147–159, 2004. 1

[11] A. Levinshtein, A. Stere, K. N. Kutulakos, D. J. Fleet, S. J. Dickinson, and K. Siddiqi. Turbopixels: Fast superpixels using geometric flows. *IEEE transactions on pattern analysis and machine intelligence*, 31(12):2290–2297, 2009. 1

[12] M.-Y. Liu, O. Tuzel, S. Ramalingam, and R. Chellappa. Entropy rate superpixel segmentation. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2097–2104. IEEE, 2011. 1, 4, 5, 6, 8

[13] Y.-J. Liu, C.-C. Yu, M.-J. Yu, and Y. He. Manifold slic: A fast method to compute content-sensitive superpixels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 651–659, 2016. 2, 4, 5, 8

[14] S. Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. 1

[15] P. Neubert and P. Protzel. Superpixel benchmark and comparison. In *Proc. Forum Bildverarbeitung*, pages 1–12, 2012. 5

[16] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *Journal of computational physics*, 79(1):12–49, 1988. 1

[17] G. Peyré, M. Péchaud, R. Keriven, and L. D. Cohen. Geodesic methods in computer vision and graphics. *Foundations and Trends® in Computer Graphics and Vision*, 5(3–4):197–397, 2010. 2

[18] X. Ren and J. Malik. Learning a classification model for segmentation. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 10–17. IEEE, 2003. 1

[19] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000. 1

[20] D. Stutz, A. Hermans, and B. Leibe. Superpixel segmentation using depth information. *RWTH Aachen University, Aachen, Germany*, 2014. 4

[21] M. Van den Bergh, X. Boix, G. Roig, B. de Capitani, and L. Van Gool. Seeds: Superpixels extracted via energy-driven sampling. In *European conference on computer vision*, pages 13–26. Springer, 2012. 1, 4, 8

[22] O. Veksler, Y. Boykov, and P. Mehrani. Superpixels and supervoxels in an energy optimization framework. In *European conference on Computer vision*, pages 211–224. Springer, 2010. 1

[23] L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE transactions on pattern analysis and machine intelligence*, 13(6):583–598, 1991. 1

[24] P. Wang, G. Zeng, R. Gan, J. Wang, and H. Zha. Structure-sensitive superpixels via geodesic distance. *International journal of computer vision*, 103(1):1–21, 2013. 2, 6

[25] S. Wang, H. Lu, F. Yang, and M.-H. Yang. Superpixel tracking. In *2011 International Conference on Computer Vision*, pages 1323–1330. IEEE, 2011. 1

[26] Z. Wang, J. Feng, S. Yan, and H. Xi. Image classification via object-aware holistic superpixel selection. *IEEE Transactions on Image Processing*, 22(11):4341–4352, 2013. 1